

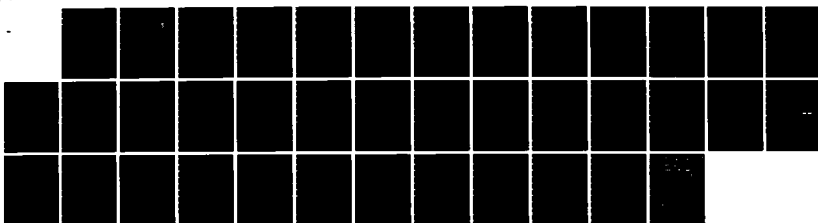
NO-A166 014

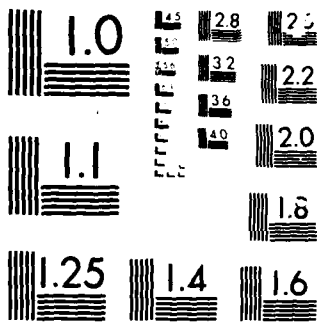
A DISTRIBUTED HYPOTHESIS-TESTING TEAM DECISION PROBLEM 1/1
WITH COMMUNICATION..(U) MASSACHUSETTS INST OF TECH
CAMBRIDGE LAB FOR INFORMATION AND D.
J PAPASTAYROU ET AL. FEB 86 LIDS-P-1530

UNCLASSIFIED

F/G 17/2

NL





MICROCOPY

CHART

AD-A166 814

February 1986

LIDS-P-1538

15

A DISTRIBUTED HYPOTHESIS-TESTING TEAM DECISION PROBLEM WITH COMMUNICATIONS COST*

DTIC
ELECTE
S
APR 17 1986
A

by
Jason Papastavrou**
Michael Athans**

This document has been approved
for public release and sale; its
distribution is unlimited.

ABSTRACT

In this paper we formulate and solve a distributed binary hypothesis-testing problem. We consider a cooperative team that consists of two decision makers (DM's); one is referred to as the primary DM and the other as the consulting DM. The team objective is to carry out binary hypothesis testing based upon uncertain measurements. The primary DM can declare his decision based only on its own measurements; however, in ambiguous situations the primary DM can ask the consulting DM for an opinion and it incurs a communications cost. Then the consulting DM transmits either a definite recommendation or pleads ignorance. The primary DM has the responsibility of making a final definitive decision. The team objective is the minimization of the probability of error, taking into account different costs for hypothesis misclassification and communication costs. Numerical results are included to demonstrate the dependence of the different decision thresholds on the problem parameters, including different perceptions of the prior information.

* Research conducted at the MIT Laboratory for Information and Decision Systems with support provided by the Office of Naval Research under contract ONR/N00014-84-k-0519 (NR 649-003)

** Room 35-406/LIDS, MIT, Cambridge, MA 02139.

This paper has been submitted to the 25th IEEE Conference on Decision and Control.

86 4 15 021

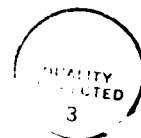
1. Introduction and Motivation.

In this paper we formulate, solve, and analyze a distributed hypothesis-testing problem which is an abstraction of a wide class of team decision problems. It represents a normative version of the "second-opinion" problem in which a primary decision maker (DM) has the option of soliciting, at a cost, the opinion of a consulting DM when faced with an ambiguous interpretation of uncertain evidence.

1.1 Motivating Examples.

Our major motivation for this research is provided by generic hypothesis-testing problems in the field of Command and Control. To be specific, consider the problem of target detection formalized as a binary hypothesis testing problem (H_0 means no target, while H_1 denotes the presense of a target). Suppose that independent noisy measurements are obtained by two geographically distributed sensors (Figure 1). One sensor, the primary DM, has final responsibility for declaring the presense or absence of a target, with different costs associated with the probability of false alarm versus the probability of missed detection. If the primary DM relied only on the measurements of his own sensor, then we have a classical centralized detection problem that has been extensively analyzed; see, for example, Van Trees [1]. If the actual measurements of the second sensor were communicated to the primary DM, we have once more a classical centralized detection problem in which we have two independent measurements on the same hypothesis; in this case, we require communication of raw data and this is expensive both from a channel bandwidth point of view and, perhaps more importantly, because radio or acoustic communication can be intercepted by the enemy.

Continuing with the target detection problem, we can arrive at the model that we shall use in the sequel by making the following assumptions which model the desire to communicate as little as possible. The primary DM can look at the data from his own sensor and attempt to arrive at a decision using a likelihood-ratio test (lrt), which yields a threshold test in the linear-Gaussian case. Quite often the primary DM can be confident about the quality of his decision. However, we can imagine that there will be instances that the data will be close to the decision threshold, corresponding to an ambiguous situation for the primary DM. In such cases it may pay off to incur a communications cost and seek some information from the other available sensor. It remains to establish what is the nature of the information to be transmitted back to the primary DM.



Decision for	
GRA&I	<input checked="" type="checkbox"/>
TAB	<input type="checkbox"/>
Unneeded	<input type="checkbox"/>
Classification	
Distribution/	
Availability Codes	
Avail and/or	
Special	

141

In our model, we assume the existence of a consulting DM having access to the data from the other sensor. We assume that the consulting DM has the ability to map the raw data from his sensor into decisions. The consulting DM is "activated" only at the request of the primary DM. It is natural to speculate that its advice will be ternary in nature: YES, I think there is a target; NO, I do not think there is a target; and, SORRY, NOT SURE MYSELF. Note that these transmitted decisions in general require less bits than the raw sensor data, hence the communication is cheap and more likely to escape enemy interception. Then, the primary DM based upon the message received from the consulting DM has the responsibility of making the final binary team decision on whether the target is present or absent.

The need for communicating with small-bit messages can be appreciated if we think of detecting an enemy submarine using passive sonar (Figure 2). We associate the primary DM with an attack submarine, and the consulting DM with a surface destroyer. Both have towed-array sonar capable of long-range enemy submarine detection. Request for information from the submarine to the destroyer can be initiated by having the sub eject a slot-buoy with a prerecorded low-power radio message. A short active sonar pulse can be used to transmit the recommendation from the destroyer to the submarine. Thus, the submarine has the choice of obtaining a "second opinion" with minimal compromise of its covert mission.

Of course, target detection is only an example of more general binary hypothesis-testing problems. Hence, one can readily extend the basic distributed team decision problem setup to other situations. For example, in the area of medical diagnosis we imagine a primary physician interpreting the outcomes of several tests. In case of doubt, he sends the patient to another consulting physician for other tests (at a dollar cost), and seeks his recommendation. However, the primary physician has the final diagnostic responsibility. Similar scenarios occur in the intelligence field where the "compartmentalization" of sensitive data, or the protection of a spy, dictate infrequent and low-bit communications. In more general military Command and Control problems, we seek insight on formalizing the need to break EMCON, and at what cost, to resolve tactical situation assessment ambiguities.

1.2 Literature Review

The solution of distributed decision problems is quite a bit different, and much more difficult, as compared to their centralized counterparts. Indeed

there is only a handful of papers that deal with solutions to distributed hypothesis-testing problems. The first attempt to illustrate the difficulties of dealing with distributed hypothesis-testing problems was published by Tenney and Sandell [2]; they point out that the decision thresholds are in general coupled. Ekchian [3] and Ekchian and Tenney [4] deal with detection networks in which downstream DM's make decisions based upon their local measurements and upstream DM decisions. Kushner and Pacut [5] introduced a delay cost (somewhat similar to the communications cost in our model) in the case that the observations have exponential distributions, and performed a simulation study. Recently, Chair and Varshney [6] have pointed out how the results in [2] can be extended in more general settings. Boettcher [7] and Boettcher and Tenney [8], [9], have shown how to modify the normative solutions in [4] to reflect human limitation constraints, and arrive in at normative/descriptive model that captures the constraints of human implementation in the presense of decision deadlines and increasing human workload; experiments using human subjects showed close agreement with the predictions of their normative/descriptive model. Finally, Tsitsiklis [10] and Tsitsiklis and Athans [11] demonstrate that such distributed hypothesis-testing problems are NP-complete; their research provides theoretical evidence regarding the inherent complexity of solving optimal distributed decision problems as compared to their centralized counterparts (which are trivially solvable).

1.3 Contributions of this Research.

The main contribution of this paper relates to the formulation and optimal solution of the team decision problem described above. Under the assumption that the measurements are conditionally independent, we show that the optimal decision rules for both the primary and the consulting DM are deterministic and are expressed as likelihood-ratio tests with constant thresholds which are tightly coupled (see Section 3 and the Appendix).

When we specialize the general results to the case that the observations are linear and the statistics are Gaussian, then we are able to derive explicit expressions for the decision thresholds for both the primary and consulting DM's (see Section 4). These threshold equations are tightly coupled, thereby necessitating an iterative solution. They provide clear-cut evidence that the DM's indeed operate as team members; their optimal thresholds are very different from those that they would use in isolation, i.e. in a non-team setting. This, of course, was the case in other versions of the distributed hypothesis-testing problem, e.g. [2].

The numerical sensitivity results (summarized in Section 5) for the linear-Gaussian case provide much needed intuitive understanding of the problem and concrete evidence that the team members operate in a more-or-less intuitive manner, especially after the fact. We study the impact of changing the communications cost and the measurement accuracy of each DM upon the decision thresholds and the overall team performance. In this manner we can obtain valuable insight on the optimal communication frequency between the DM's. As to be expected, as the communication cost increases, the frequency of communication (and asking for a second opinion) decreases, and the team performance approaches that of the primary DM operating in isolation. In addition, we compare the overall distributed team performance to the centralized version of the problem in which the primary DM had access, at no cost, to both sets of observations. In this manner, we can study the degree of inherent performance degradation to be expected as a consequence of enforcing the distributed decision architecture in the overall decision making process.

Finally, we study the team performance degradation when one of the team members, either the primary or the consulting DM, has an erroneous estimate of the hypotheses prior probabilities. This corresponds to mildly different mental models of the prior situation assesment; see Athans [12]. As expected the team performance is much more sensitive to misperceptions by the primary DM as compared to similar misperceptions by the consulting DM. This implies that, if team training reduces misperceptions on the part of the DM's, the greatest payoff is obtained in training the primary DM.

2. Problem Definition

The problem is one of hypothesis testing. The team has to choose among two alternative hypotheses H_0 and H_1 , with a priori probabilities

$$P(H_0)=p_0 \qquad P(H_1)=p_1 \qquad (1)$$

Each of two DM's, one called primary (DM A) and one consulting (DM B), receives an uncertain measurement y_a and y_b respectively (Figure 1), distributed with known joint probability density functions

$$P(y_\alpha, y_\beta | H_i) \quad ; \quad i=0,1. \quad (2)$$

The final decision of the team u_r (0 or 1, indicating H_0 or H_1 to be true) is the responsibility of the primary DM. DM A initially makes a preliminary decision u_α where it can either decide (0 or 1) on the basis of its own data (ie y_α), or at a cost ($C \geq 0$) can solicit DM B's opinion ($u_\beta=1$), prior to making the commital decision.

The consulting DM's decision u_β consists of three distinct messages (call them : x, y and z) and is activated only when asked. We decided to assign three messages to DM B, because we wanted to have one message indicating each of the two hypotheses and one message indicating that the consulting DM is 'not sure.' In fact, we proved that the optimal content for the messages of DM B is the one mentioned above.

When the message from DM B is received, the burden shifts back to the primary DM, which is called to make the commital decision of the team based on his own data and the information from the consulting DM.

We now define the following cost function :

$$J : \{0,1\} \times \{H_0, H_1\} \rightarrow R \quad (3)$$

with $J(u_r, H_i)$ being the cost incurred by the team choosing u_r , when H_i is true.

Then, the optimality criterion for the team is a function

$$J^* : \{0,1,1\} \times \{0,1\} \times \{H_0, H_1\} \rightarrow R \quad (4)$$

with :

$$J^*(u_\alpha, u_\beta, H_i) = \begin{cases} J(u_r, H_i) + C & ; \quad u_\alpha=1 \text{ (information requested)} \\ J(u_r, H_i) & ; \quad \text{otherwise} \end{cases} \quad (5)$$

The cost structure of the problem is the usual cost structure used in hypothesis testing problems, but also includes the non-negative communication cost, which the team 'pays' when the DM A decides to obtain the consulting DM's information.

Remark : According to the rules of the problem, when the preliminary decision u_α of the primary DM is 0 or 1, then the final team decision is 0 or 1 respectively (ie $P(u_r=1 | u_\alpha=1)=1$ for $i=0,1$).

The objective of the decision strategies will be to minimize the expected cost incurred

$$\min E[J^*(u_\alpha, u_\gamma, H)] \quad (6)$$

where the minimization is over the decision rules of the two DMs. Note that the decision rule of the consulting DM is implicitly included in the cost function, through the final team decision u_γ (which is a function of the decision of the consulting DM).

All the prior information is known to both DMs. The only information they do not share is their observations. Each DM knows only its own observation and, because of the conditional independence assumption, nothing about the other DM's observation.

The problem can now be stated as follows :

Problem : Given p_0, p_1 , the distributions $P(y_\alpha, y_\beta | H_i)$ for $i=0,1$ with $y_\alpha \in Y_\alpha$, $y_\beta \in Y_\beta$, and the cost function J^* , find the decision rules u_α, u_β and u_γ as functions

$$u_\alpha : Y_\alpha \rightarrow \{0,1\} \quad (7)$$

$$u_\beta : Y_\beta \rightarrow \{x,y,z\} \quad (8)$$

and

$$u_\gamma : Y_\alpha \times \{x,y,z\} \rightarrow \{0,1\} \quad (9)$$

(subject to : $P(u_\gamma=1 | u_\alpha=1)=1$ for $i=0,1$), which minimize the expected cost.

NOTE : The centralized counterpart of the problem, where a single DM receives both observations is a well known problem. The solution is deterministic and given by a likelihood ratio test (lrt). That is :

$$u_0 : Y_\alpha \times Y_\beta \rightarrow \{0,1\} \quad (10)$$

with

$$u_0(y_\alpha, y_\beta) = \begin{cases} 0 & ; \quad \Lambda(y_\alpha, y_\beta) \geq t \\ 1 & ; \quad \text{otherwise} \end{cases} \quad (11)$$

where

$$\begin{aligned}\Delta(y_\alpha, y_\beta) &= [P(y_\alpha, y_\beta | H_0)p_0] / [P(y_\alpha, y_\beta | H_1)p_1] \\ &= P(H_0 | y_\alpha, y_\beta) / P(H_1 | y_\alpha, y_\beta)\end{aligned}\quad (12)$$

and t is a precomputed threshold

$$t = [J(0, H_1) - J(1, H_1)] / [J(1, H_0) - J(0, H_0)] \quad (13)$$

provided $J(1, H_0) > J(0, H_0)$. Thus, the difficulty of our problem arises because of its decentralized nature.

We will show that, under certain assumptions, the most restrictive of which is conditional independence of the observations, the optimal decision rules for the *Problem* are deterministic and given by lrt's with constant thresholds. The thresholds of the two DMs are coupled, indicating that the DMs work as a team rather than individuals.

3. About the Solution to the General Problem

In order to be able to solve the *Problem*, we make the following assumptions.

$$\text{ASSUMPTION 1: } J(1, H_0) > J(0, H_0) ; J(0, H_1) > J(1, H_1) \quad (14)$$

or it is more costly for the team to err than to be correct.

This logical assumption is made in order to motivate the team members to avoid erring and in order to enable us to algebraically put the optimal decisions in lrt form.

$$\text{ASSUMPTION 2: } P(y_\alpha | y_\beta, H_i) = P(y_\alpha | H_i) ; P(y_\beta | y_\alpha, H_i) = P(y_\beta | H_i) ; i=0,1 \quad (15)$$

or the observations y_α and y_β are conditionally independent.

This assumption removes the dependence of the one observation on the other and thus allows us to write the optimal decision rules as lrt's with constant thresholds.

ASSUMPTION 3 : Without loss of generality assume that :

$$\frac{P(u_p=x | u_\alpha=I, H_0)}{P(u_p=x | u_\alpha=I, H_1)} \geq \frac{P(u_p=v | u_\alpha=I, H_0)}{P(u_p=v | u_\alpha=I, H_1)} \geq \frac{P(u_p=z | u_\alpha=I, H_0)}{P(u_p=z | u_\alpha=I, H_1)} \quad (16)$$

This assumption is made in order to distinguish between the messages of DM B.

As shown in detail in the Appendix, the optimal decision rules for all three decisions of our problem (u_α , u_p , u_r) are given by *deterministic* functions which are expressed as *likelihood ratio tests* with *constant* thresholds. The three thresholds of the primary DM (two for u_α and one for u_r) and the two thresholds of the consulting DM (for u_p) can not be obtained in closed form. They are coupled, that is the thresholds of one DM are given as functions of the thresholds of the other DM.

Another important result is that, when the optimal decision rules are employed and the consulting DM's decision is x (or z), then the optimal final decision rule of the primary DM is always 0 (or 1) :

$$P(u_r=0 | u_\alpha=I, u_p=x, y_\alpha)=1 \text{ for all } y_\alpha \in \{y_\alpha | P(u_\alpha=I | y_\alpha)=1, y_\alpha \in Y_\alpha\} \quad (17)$$

and

$$P(u_r=1 | u_\alpha=I, u_p=z, y_\alpha)=1 \text{ for all } y_\alpha \in \{y_\alpha | P(u_\alpha=I | y_\alpha)=1, y_\alpha \in Y_\alpha\} \quad (18)$$

Thus, we can simplify our notation by changing the DM B decisions from x to 0, from z to 1 and from v to ? (which is interpreted as "I am not sure"). The team's decision process can be now described as follows : Each of the two DMs receives an observation. Then, the primary DM can either make the final decision (0 or 1) or can decide to incur the communication cost ($u_\alpha=I$) and pass the responsibility of the final decision to the consulting DM. When called upon, the consulting DM can either make the final decision or shift the burden back to DM A ($u_p=?$), in which case the primary DM is forced to make the final decision, based on its own observation (y_α) and the fact that DM B decided $u_p=?$.

A detailed presentation of the facts discussed above can be found in the Appendix.

4. A Gaussian Example

We now present detailed threshold equations for the case where the probability distributions of the two observations are Gaussian. We selected the Gaussian distribution, despite its cumbersome algebraic formulae, because of its generality. Our objective is to perform numerical sensitivity analysis to the solution of this example, in order to gain information on the team 'activities.'

We assume that the observations are distributed with the following Gaussian distributions :

$$y_{\alpha} \sim N(\mu, \sigma_{\alpha}^2) \quad ; \quad y_{\beta} \sim N(\mu, \sigma_{\beta}^2) \quad (19)$$

The two alternative hypotheses are :

$$H_0 : \mu = \mu_0 \quad \text{or} \quad H_1 : \mu = \mu_1 \quad (20)$$

Without loss of generality, assume that :

$$\mu_0 < \mu_1 \quad (21)$$

The rest of the notation is the same as in the general problem described above.

We can show that the optimum decision rules for this example are given by thresholds on the *observation* axes, as shown in Figure 3. Before presenting the equations of the thresholds, we define some variables.

Ψ_{α}^l : lower threshold of DM A	Ψ_{α}^u : upper threshold of DM A
Ψ_{α}^f : threshold for the final decision of DM A	
Ψ_{β}^l : lower threshold of DM B	Ψ_{β}^u : upper threshold of DM B

$$\Phi_i^j(k) = \int_{-\infty}^{\frac{\Psi_i^j - \mu_k}{\sigma_i}} (2\pi)^{-0.5} \exp(-0.5 x^2) dx \quad \text{for } i=\alpha, \beta \ ; \ j=l, f, u \ ; \ k=0, 1$$

Note that the above function is the well-known error function, presented with notational modifications to fit the purposes of the problem.

$$W^1 = 0.5 [\Phi_\beta^u(0) - \Phi_\beta^u(1)] \quad (22)$$

$$W^2 = \frac{\Phi_\beta^u(0) - \Phi_\beta^l(0) - \Phi_\beta^u(1) + \Phi_\beta^l(1) + \Phi_\beta^l(0)\Phi_\beta^u(1) - \Phi_\beta^u(0)\Phi_\beta^l(1)}{\Phi_\beta^u(0) - \Phi_\beta^l(0) + \Phi_\beta^u(1) - \Phi_\beta^l(1)} \quad (23)$$

$$W^3 = \frac{\Phi_\beta^l(0)\Phi_\beta^u(1) - \Phi_\beta^l(1)\Phi_\beta^u(0)}{\Phi_\beta^u(0) - \Phi_\beta^l(0) + \Phi_\beta^u(1) - \Phi_\beta^l(1)} \quad (24)$$

$$W^4 = 0.5 [\Phi_\beta^u(0) - \Phi_\beta^u(1)] \quad (25)$$

$$\psi_\alpha^* = (\mu_0 + \mu_1)/2 + [\sigma_\alpha^2 / (\mu_1 - \mu_0)] \ln[p_0 / (1 - p_0)] \quad (26)$$

$$\psi_\beta^* = (\mu_0 + \mu_1)/2 + [\sigma_\beta^2 / (\mu_1 - \mu_0)] \ln[p_0 / (1 - p_0)] \quad (27)$$

In (26) and (27), the (centralized) maximum likelihood estimators for each DM are defined.

COROLLARY 1 : If $P(u_\alpha=1) > 0$ (i.e. information is requested for some y_α) and if $P(u_\beta=? | u_\alpha=1) > 0$ (i.e. "I am not sure" is returned for some y_β , when information is requested), then the optimal final decision rule of the primary DM is a deterministic function defined by :

$$\gamma_f(y_\alpha) = \begin{cases} 0 & \text{if } y_\alpha \leq \psi_\alpha^f \\ 1 & \text{if } y_\alpha > \psi_\alpha^f \end{cases} \quad (28)$$

where :

$$\psi_\alpha^f = \psi_\alpha^* + \frac{\sigma_\alpha^2}{\mu_1 - \mu_0} \ln \left(\frac{\Phi_\beta^u(0) - \Phi_\beta^l(0)}{\Phi_\beta^u(1) - \Phi_\beta^l(1)} \right) \quad (29)$$

COROLLARY 2 : If $P(u_\alpha=1) > 0$ (i.e. information is requested for some y_α) and the primary DM's final decision rule is the one given by Corollary 1, then the optimal decision rule of the consulting DM is a deterministic function defined by :

$$\gamma_\beta(y_\beta) = \begin{cases} 0 & \text{if } y_\beta \leq \psi_\beta^l \\ ? & \text{if } \psi_\beta^l < y_\beta \leq \psi_\beta^u \\ 1 & \text{if } \psi_\beta^u < y_\beta \end{cases} \quad (30)$$

where :

$$\psi_\beta^1 = \min \left\{ \psi_\beta^* + \frac{\sigma_\beta^2}{\mu_1 - \mu_0} \ln \left(\frac{\Phi_\alpha^u(0) - \Phi_\alpha^f(0)}{\Phi_\alpha^u(1) - \Phi_\alpha^f(1)} \right), \psi_\beta^* + \frac{\sigma_\beta^2}{\mu_1 - \mu_0} \ln \left(\frac{\Phi_\alpha^u(0) - \Phi_\alpha^1(0)}{\Phi_\alpha^u(1) - \Phi_\alpha^1(1)} \right) \right\} \quad (31)$$

and :

$$\psi_\beta^u = \max \left\{ \psi_\beta^* + \frac{\sigma_\beta^2}{\mu_1 - \mu_0} \ln \left(\frac{\Phi_\alpha^u(0) - \Phi_\alpha^1(0)}{\Phi_\alpha^u(1) - \Phi_\alpha^1(1)} \right), \psi_\beta^* + \frac{\sigma_\beta^2}{\mu_1 - \mu_0} \ln \left(\frac{\Phi_\alpha^f(0) - \Phi_\alpha^1(0)}{\Phi_\alpha^f(1) - \Phi_\alpha^1(1)} \right) \right\} \quad (32)$$

COROLLARY 3 : Given that the final decision rule employed by the primary DM is the one of Corollary 1 and that the decision rule employed by the consulting DM is the one of Corollary 2, then the optimal decision rule for the preliminary decision u_α of the primary DM is a deterministic function defined by :

$$\gamma_\alpha(u_\alpha) = \begin{cases} 0 & \text{if } u_\alpha \leq \psi_\alpha^1 \\ 1 & \text{if } \psi_\alpha^1 < u_\alpha \leq \psi_\alpha^u \\ 1 & \text{if } \psi_\alpha^u < u_\alpha \end{cases} \quad (33)$$

where :

$$\psi_\alpha^1 = \begin{cases} \psi_\alpha^* + \frac{\sigma_\alpha^2}{\mu_1 - \mu_0} \ln \left(\frac{1 - \Phi_\beta^1(0) + C}{1 - \Phi_\beta^1(1) - C} \right) & ; \quad 0 \leq C < \min \{W^1, W^2\} \\ \psi_\alpha^* + \frac{\sigma_\alpha^2}{\mu_1 - \mu_0} \ln \left(\frac{1 - \Phi_\beta^u(0) + C}{1 - \Phi_\beta^u(1) - C} \right) & ; \quad W^2 < C \leq W^4 \\ \psi_\alpha^* & ; \quad \text{otherwise} \end{cases} \quad (34)$$

and :

$$\psi_\alpha^u = \begin{cases} \psi_\alpha^* + \frac{\sigma_\alpha^2}{\mu_1 - \mu_0} \ln \left(\frac{\Phi_\beta^1(0) - C}{\Phi_\beta^1(1) + C} \right) & ; \quad 0 \leq C < \min \{W^3, W^4\} \\ \psi_\alpha^* + \frac{\sigma_\alpha^2}{\mu_1 - \mu_0} \ln \left(\frac{\Phi_\beta^u(0) - C}{\Phi_\beta^u(1) + C} \right) & ; \quad W^3 < C \leq W^1 \\ \psi_\alpha^* & ; \quad \text{otherwise} \end{cases} \quad (35)$$

REMARK : Observe that the equations of all the thresholds include (and possibly reduce to) a "centralized" part (ψ_i^*) indicating the relation of our problem to its centralized counterpart.

5. NUMERICAL SENSITIVITY ANALYSIS

We now perform sensitivity analysis to the solution of the Gaussian example. Our objective is to analyze the effects on the team performance from varying the parameters of our problem, in order to obtain better understanding of the decentralized team decision mechanism. We vary the quality of the observations of each DM (the variance of each DM), the a priori likelihood of the hypotheses and the communication cost. Finally, we study the effects of different a priori knowledge for each DM.

We use the following 'minimum error' cost function :

$$J(u_i, H_i) = \begin{cases} 0 & ; \quad u_i = 1 \\ 1 & ; \quad u_i \neq 1 \end{cases} \quad (36)$$

We do not need to vary the cost function, because this would be mathematically equivalent to varying the a priori probabilities of the two hypotheses.

5.1 Effects of varying the quality of the observations of the Primary DM

Denote :

$C1^*$ = cost incurred if the consulting DM makes the decision alone .

We distinguish two cases depending on the cost associated with the information (ie the of quality of information)

CASE 1: $\min(P_0, 1 - P_0) \leq C1^* + C$

As the variance of the primary DM increases, it becomes less costly for the team to have the primary DM always decide the more likely hypothesis, then request for information. This occurs because the observation of DM A becomes increasingly worthless. Thus, the primary DM progressively ignores its observation and in order to minimize cost has to choose

between "de facto" deciding the more likely hypothesis (and incurring cost equal to the probability of the least likely hypothesis) or "de facto" requesting for information (and thus incurring the communication cost plus the cost of the consulting DM). In this case, the prior is less than the latter and so the optimum decision of the primary DM, as its variance tends to infinity is to always decide the more likely hypothesis (Figure 4, $P_0 = .8$). Thus :

$$\lim_{\sigma_0^2 \rightarrow \infty} P(u_\alpha = 1) \rightarrow 0$$

Moreover, the percentage gain in cost achieved by the team of DMs, relative to the cost incurred by a single DM obtaining a single observation, asymptotically goes to 0, as the variance of the primary DM goes to infinity (Figure 5, $P_0 = .8$).

An interesting insight can be obtained from Figure 5 ($P_0 = .8$). As the variance of the primary DM increases the percentage improvement in cost (defined above) is initially increasing and then decreasing asymptotically to zero. The reason for this is that for very small variances, the observations of the primary DM are so good that it does not need the information of the consulting DM. As the variance increases, the primary DM makes better use of the information and so the percentage improvement increases. But, at a certain point as the quality of the observations worsens, the primary DM finds less costly to start declaring more often the more likely hypothesis (ie to bias its decision towards the more likely hypothesis) than requesting for information, for reasons mentioned above, and so the percentage improvement from then on decreases.

CASE 2: $\min(P_0, 1 - P_0) > C1^* + C$

With reasoning similar to the above, we obtain that (Figure 4, $P_0 = .5$) :

$$\lim_{\sigma_0^2 \rightarrow \infty} P(u_\alpha = 1) \rightarrow 1$$

Moreover, the percentage improvement is strictly increasing (and keeps increasing to a precomputable limit ; Figure 5, $P_0 = .5$). This reinforces the last point we made in (Case 1) above. Since in the present case it is always less costly for the primary DM to request and use the information than to bias its decision towards the more likely hypothesis, the percentage improvement curve does not exhibit the non-monotonic behavior observed in (Case 1) above (where $P_0 = .8$).

5.2 Effects of varying the quality of the observations of the Consulting DM

As the variance of the consulting DM's observations increases, less information is requested by the primary DM, that is the primary DM's upper and lower thresholds move closer to each other (Figure 6). This is something we expected, since information of lesser quality is less profitable (more costly) to the team of DMs.

We should note here that the thresholds of a DM is an alternative way of representing the probabilities of the DM's decisions, since the decision regions are characterized by the thresholds. For example :

$$P(u_{\alpha}=1) = \sum \int_H P(y_{\alpha} | H) P(H) dy_{\alpha} : \psi_{\alpha}^l < y_{\alpha} < \psi_{\alpha}^u$$

The thresholds of the consulting DM demonstrate some interesting points of the team behavior (Figure 7). For small values of the variance they are very close together, as the quality of the observations is very good and so the consulting DM is willing to make the final team decision. As the variance increases, DM B becomes more willing to return $u_p=?$ (i.e. "I am not sure") and let DM A make the final team decision. As the variance continues to increase, the thresholds of the consulting DM converge again. It might seem counter-intuitive, but there is a simple explanation. The consulting DM recognizes that the primary DM, despite knowing that the quality of the consulting DM's information is bad, is willing to incur the communication cost to obtain the information. This indicates that the primary DM is 'confused', that is, the a posteriori probabilities of the two hypotheses (given its observation) are very close together. Hence, the consulting DM becomes more willing to make the final decision. After a certain point ($\sigma_p^2 \approx 62.4$) the primary DM does find it worthwhile to request for information at all.

REMARK : Note in Figure 7 that the thresholds of the consulting DM converge to 1 which would have been the maximum likelihood threshold had the a priori probabilities of the two hypotheses been equal. But, the a priori probabilities *which the consulting DM uses in its calculations* are functions of the given a priori probabilities (ie p_j) and the fact that the primary DM requested for information (ie $P(u_{\alpha}=1 | H_j)$). In fact, the consulting DM uses as its a priori probabilities its own estimates of the

primary DM's a posteriori probabilities. That is :

$$P(H_0 | u_\alpha = I) = \frac{P(H_0) P(u_\alpha = I | H_0)}{\sum_H P(H) P(u_\alpha = I | H)} \quad (37)$$

From the above, we deduce that for large variances ($\sigma_p^2 \approx 62$) the estimates, of the consulting DM, for the a posteriori probabilities of the primary DM are very close to .5, reenforcing our point about the primary DM "being confused."

Finally, it is clear, that as the variance of the consulting DM increases, the percentage gain in cost, achieved by the team of DMs, decreases to 0, since the primary DM eventually makes all the decisions alone (centralized).

5.3 Effects of varying the Communication Cost

Increasing the communication cost is very similar to increasing the variance of the consulting DM, since in both cases the team "gets less for its money" (because the team has to incur an increased cost, either in the form of an increased communication cost, or in the form of the final cost, because of the worse performance of the consulting DM).

The thresholds of the primary DM, exhibit the same behavior as in 5.2 above (converging together at $C \approx .35$). The thresholds of the consulting DM (Figure 8) converge together for the same reasons as in 5.2 above. Of course, the thresholds do not start together for small values of the communication cost (as in 5.2), because low communication cost does not imply ability for the consulting DM to make accurate decisions. In fact, for small values of the communication cost, DM A is compelled to request for information more often than what is really needed and so the consulting DM returns more often $u_p = ?$ (ie "I am not sure") and lets DM A make the team final decision.

Again it is clear that, as the communication cost increases, the percentage gain achieved by the team of the DMs decreases to zero (as the communication becomes more costly and less frequent, until we reach the centralized case).

5.4 Effects of varying the a priori probabilities of the hypotheses

This case does not present many interesting points. As expected, there is

symmetry in the performance of the team around the line $p_0 = 0.5$. The closer p_0 is to 0.5 the more often information is requested by DM A (Figure 9) and the more often "I am not sure" is returned by DM B (Figure 10). This is understandable, because the closer p_0 is to 0.5, the bigger the a priori uncertainty. Consequently, the percentage improvement achieved by the team of the DMs is monotonically increasing with p_0 from 0 to 0.5 and monotonically decreasing from 0.5 to 1.

5.5 Effects of imperfect a priori information

CASE 1: Only the consulting DM knows the true p_0

From Figure 11, where the true p_0 is 0.8, we deduce that our model is relatively robust. If the primary DM's erroneous p_0 is anywhere between 0.7 and 0.9, performance of the team will be not more than 10% away from the optimum.

CASE 2: Only the primary DM knows the true p_0

As we see in Figure 12, where the true p_0 is 0.8, our model exhibits remarkable robustness qualities. If the consulting DM's erroneous p_0 is as far out as 0.01, the performance of the team will not be further than 7% away from the optimal. This can be explained by looking at the consulting DM's thresholds as functions of p_0 (Figure 13). We observe that for values of p_0 between 0.01 and 0.99, the thresholds do not change by much. This occurs because, as explained in detail in 5.2 above, the consulting DM knows that the primary DM requests for information when its a posteriori probabilities of the two hypotheses are roughly equal, which is the case indeed. As already stated, the consulting DM uses as its a priori probabilities its estimates of the a posteriori probabilities of the primary DM. Therefore, the consulting DM's estimates of the primary DM's a posteriori probabilities are good, besides the discrepancy in p_0 , and the team's performance is not tempered by much.

Acknowledgment

The authors would like to thank Professor John N. Tsitsiklis for his valuable suggestions.

6. REFERENCES

- [1] Van Trees, H.L. (1969)
Detection, Estimation and Modulation Theory, vol. I.
New York : Wiley, 1969
- [2] Tenney, R.R., and Sandell, N.R. (1981)
Detection with Distributed Sensors
IEEE Trans. of Aerospace and Electronic Systems
17, 4 (July 1981), pp. 501-509
- [3] Ekchian, L.K. (1982)
Optimal Design of Distributed Detection Networks
Ph.D. Dissertation, Dept. of Elect. Eng. and Computer Science
Mass. Inst. of Technology
- [4] Ekchian, L.K., and Tenney, R.R. (1982)
Detection Networks
Proceedings of the 21st IEEE Conference on Decision and Control
pp. 686-691
- [5] Kushner, H.J., and Pacut, A. (1982)
A Simulation Study of a Decentralized Detection Problem
IEEE Trans. on Automatic Control, 27, 5 (Oct. 1982)
pp. 1116-1119
- [6] Chair, Z., and Varshney, P.K. (1986)
Optimal Data Fusion in Multiple Sensor Detection Systems
IEEE Trans. on Aerospace and Electronic Systems
21, 1, (January 1986), pp. 98-101
- [7] Boettcher, K.L. (1985)
A Methodology for the Analysis and Design of Human
Information Processing Organizations
Ph.D. Dissertation, Dept. of Elect. Eng. and Computer Science
Mass. Inst. of Technology
- [8] Boettcher, K.L., and Tenney, R.R. (1985)
On the Analysis and Design of Human Information Processing
Organizations
Proceedings of the 8th MIT/ONR Workshop on C³ Systems
pp. 69-74
- [9] Boettcher, K.L., and Tenney, R.R. (1985)
Distributed Decisionmaking with Constrained Decision Makers
A Case Study
Proceedings of the 8th MIT/ONR Workshop on C³ Systems
pp. 75-79

- [10] Tsitsiklis, J.N. (1984)
 Problems in Decentralized Decision Making and Computation
 Ph.D. Dissertation, Dept. of Elec. Eng. and Computer Science
 Mass. Inst. of Technology
- [11] Tsitsiklis, J.N., and Athans, M. (1985)
 On the Complexity of Decentralized Decision Making
 and Detection Problems
 IEEE Trans. on Automatic Control, 30, 5, (May 1985)
 pp. 440-446

APPENDIX : Solution to the General Problem

ASSUMPTION 1 : $J(1, H_0) > J(0, H_0)$; $J(0, H_1) > J(1, H_1)$ (38)
 or it is more costly for the team to err than to be correct.

This logical assumption is made in order to motivate the team members to avoid erring and in order to enable us to put the optimal decisions in 1st form.

ASSUMPTION 2 : $P(y_\alpha | y_\beta, H_i) = P(y_\alpha | H_i)$; $P(y_\beta | y_\alpha, H_i) = P(y_\beta | H_i)$; $i=0,1$ (39)
 or the observations y_α and y_β are conditionally independent.

This assumption removes the dependence of the one observation on the other and thus allows us, as we are about to show, to write the optimal decision rules as 1st's with constant thresholds.

ASSUMPTION 3 : Without loss of generality assume that :

$$\frac{P(u_\beta = x | u_\alpha = 1, H_0)}{P(u_\beta = x | u_\alpha = 1, H_1)} \geq \frac{P(u_\beta = v | u_\alpha = 1, H_0)}{P(u_\beta = v | u_\alpha = 1, H_1)} \geq \frac{P(u_\beta = z | u_\alpha = 1, H_0)}{P(u_\beta = z | u_\alpha = 1, H_1)} \quad (40)$$

This assumption is made in order to be able to distinguish between the messages of DM B.

THEOREM 1 : Given decision rules u_α and u_β , and that information is requested by the primary DM for some $y_\alpha \in Y_\alpha$ (i.e. $P(u_\alpha=1) > 0$), then the optimal final decision of the primary DM after the information has been received, can be expressed as a deterministic function

$$\gamma_f : Y_\alpha \times \{x, v, z\} \rightarrow \{0, 1\}$$

which is defined as likelihood ratio tests :

$$\gamma_f(y_\alpha, u_\beta) = \begin{cases} 0 & ; \text{ if } u_\beta=1 \text{ and } \Lambda_\alpha(y_\alpha) \geq \alpha_i \\ 1 & ; \text{ otherwise} \end{cases} \quad \text{for } i=x, v, z \quad (41)$$

where :

$$\Lambda_\alpha(y_\alpha) = \frac{p_0 P(y_\alpha | H_0)}{p_1 P(y_\alpha | H_1)} \quad (42)$$

and :

$$\alpha_i = \frac{P(u_\beta=1 | u_\alpha=1, H_1) [J(1, H_1) - J(0, H_1)]}{P(u_\beta=1 | u_\alpha=1, H_0) [J(1, H_0) - J(0, H_0)]} \quad ; \quad i=x, v, z \quad (43)$$

Remark : (29) is the equation for the corresponding threshold of the Gaussian example.

THEOREM 2 : Given the optimal decision rule u_f (derived in Theorem 1), a decision rule u_α and that information is requested for some $y_\alpha \in Y_\alpha$ (i.e. $P(u_\alpha=1) > 0$), the optimal decision rule of the consulting DM is a deterministic function

$$\gamma_\beta : Y_\beta \rightarrow \{x, v, z\}$$

defined as the following likelihood ratio tests :

$$\gamma_\beta(y_\beta) = \begin{cases} x & \text{if } \Lambda_\beta(y_\beta) \geq b_1 \text{ and } \Lambda_\beta(y_\beta) \geq b_2 \\ v & \text{if } \Lambda_\beta(y_\beta) < b_1 \text{ and } \Lambda_\beta(y_\beta) \geq b_3 \\ z & \text{if } \Lambda_\beta(y_\beta) < b_2 \text{ and } \Lambda_\beta(y_\beta) < b_3 \end{cases} \quad (44)$$

where :

$$\Delta_p(y_p) = \frac{p_0 P(y_p | H_0)}{p_1 P(y_p | H_1)} \quad (45)$$

and :

$$b_1 = \frac{P(u_\alpha = I | H_1) \sum_{u_r} [P(u_r | u_\alpha = I, u_p = v, H_1) - P(u_r | u_\alpha = I, u_p = x, H_1)] J(u_r, H_1)}{P(u_\alpha = I | H_0) \sum_{u_r} [P(u_r | u_\alpha = I, u_p = x, H_0) - P(u_r | u_\alpha = I, u_p = v, H_0)] J(u_r, H_0)} \quad (46)$$

$$b_2 = \frac{P(u_\alpha = I | H_1) \sum_{u_r} [P(u_r | u_\alpha = I, u_p = z, H_1) - P(u_r | u_\alpha = I, u_p = x, H_1)] J(u_r, H_1)}{P(u_\alpha = I | H_0) \sum_{u_r} [P(u_r | u_\alpha = I, u_p = x, H_0) - P(u_r | u_\alpha = I, u_p = z, H_0)] J(u_r, H_0)} \quad (47)$$

$$b_3 = \frac{P(u_\alpha = I | H_1) \sum_{u_r} [P(u_r | u_\alpha = I, u_p = z, H_1) - P(u_r | u_\alpha = I, u_p = v, H_1)] J(u_r, H_1)}{P(u_\alpha = I | H_0) \sum_{u_r} [P(u_r | u_\alpha = I, u_p = v, H_0) - P(u_r | u_\alpha = I, u_p = z, H_0)] J(u_r, H_0)} \quad (48)$$

Equivalently, we can write :

$$y_p(y_p) = \begin{cases} x & \text{if } \Delta_p(y_p) \geq \beta_1 \\ v & \text{if } \Delta_p(y_p) < \beta_1 \text{ and } \Delta_p(y_p) \geq \beta_2 \\ z & \text{if } \Delta_p(y_p) < \beta_2 \end{cases} \quad (49)$$

where :

$$\beta_1 = \max \{ b_1, b_2 \} \quad (50)$$

and :

$$\beta_2 = \min \{ b_2, b_3 \}. \quad (51)$$

Remark : (31) and (32) are the equations of the corresponding thresholds of the Gaussian example.

LEMMA 1 : Given the decision rule u_β of the consulting DM and the final decision rule u_γ of the primary DM, the preliminary decision rule u_α of the primary DM can be expressed as a deterministic function

$$\gamma_\alpha : Y_\alpha \rightarrow \{0, 1, I\}$$

defined as the following *degenerate* (because the thresholds are functions of y_α) Irt's :

$$\gamma_\alpha(y_\alpha) = \begin{cases} 0 & \text{if } \Delta_\alpha(y_\alpha) \geq \theta_1 \text{ and } \Delta_\alpha(y_\alpha) \geq \theta_2 \\ 1 & \text{if } \Delta_\alpha(y_\alpha) < \theta_2 \text{ and } 1/\Delta_\alpha(y_\alpha) < 1/\theta_3 \\ I & \text{if } \Delta_\alpha(y_\alpha) < \theta_1 \text{ and } 1/\Delta_\alpha(y_\alpha) \geq 1/\theta_3 \end{cases} \quad (52)$$

where $\Delta_\alpha(y_\alpha)$ is defined in (42) and :

$$\theta_1 = \frac{J(0, H_1) - J(1, H_1)}{J(1, H_0) - J(0, H_0)} \quad (53)$$

$$\theta_2 = \frac{\sum_{u_\gamma, u_\beta} P(u_\gamma | u_\alpha = I, u_\beta, y_\alpha) P(u_\beta | u_\alpha = I, H_1) [J(u_\gamma, H_1) + C] - J(0, H_1)}{J(0, H_0) - \sum_{u_\gamma, u_\beta} P(u_\gamma | u_\alpha = I, u_\beta, y_\alpha) P(u_\beta | u_\alpha = I, H_0) [J(u_\gamma, H_0) + C]} \quad (54)$$

$$\theta_3 = \frac{\sum_{u_\gamma, u_\beta} P(u_\gamma | u_\alpha = I, u_\beta, y_\alpha) P(u_\beta | u_\alpha = I, H_1) [J(u_\gamma, H_1) + C] - J(1, H_1)}{J(1, H_0) - \sum_{u_\gamma, u_\beta} P(u_\gamma | u_\alpha = I, u_\beta, y_\alpha) P(u_\beta | u_\alpha = I, H_0) [J(u_\gamma, H_0) + C]} \quad (55)$$

We proceed to show that the thresholds derived above are *independent* of y_α .

COROLLARY 1 : If for some y_α information is requested, according to the rule of Lemma 1 and $u_\beta = x$ (or z) is returned, then the optimal final decision u_γ of the primary DM is always 0 (or 1); that is :

$$P(u_r=0 | u_\alpha=1, u_\beta=x, y_\alpha) = 1 \quad \text{for all } y_\alpha \in \{y_\alpha | P(u_\alpha=1 | y_\alpha)=1, y_\alpha \in Y_\alpha\} \quad (56)$$

and

$$P(u_r=1 | u_\alpha=1, u_\beta=z, y_\alpha) = 1 \quad \text{for all } y_\alpha \in \{y_\alpha | P(u_\alpha=1 | y_\alpha)=1, y_\alpha \in Y_\alpha\} \quad (57)$$

Remark : From Corrolary 1 we can now give another interpretation to the team procedure: the primary DM can decide 0 or 1 using his own observation or can decide, because of uncertainty, to incur the communication cost (C) and shift the burden of the decision to the consulting DM. Then it is the consulting DM's turn to choose between deciding 0 or 1, or, because of uncertainty, shifting the burden back (at no cost) to the primary DM, which is required to make the final decision given his observation and the fact that the consulting DM's observation is not good enough for the consulting DM to make the final decision.

According to the above, we can simplify our notation of the consulting DM's messages by changing x to 0, z to 1 and v to ? (which is interpreted as the consulting DM saying "I am not sure").

Define the following secondary variables :

$$\Delta J_0 = J(1, H_0) - J(0, H_0) \quad (58)$$

$$\Delta J_1 = J(0, H_1) - J(1, H_1) \quad (59)$$

$$W^1 = \frac{\Delta J_0 \Delta J_1 [P(u_\beta=1 | H_1) - P(u_\beta=1 | H_0)]}{\Delta J_0 + \Delta J_1} \quad (60)$$

$$W^2 = \frac{\Delta J_0 \Delta J_1 [P(u_\beta=? | H_0)P(u_\beta=1 | H_1) - P(u_\beta=? | H_1)P(u_\beta=1 | H_0)]}{\Delta J_0 P(u_\beta=? | H_0) + \Delta J_1 P(u_\beta=? | H_1)} \quad (61)$$

$$W^3 = \frac{\Delta J_0 \Delta J_1 [P(u_\beta=? | H_1)P(u_\beta=0 | H_0) - P(u_\beta=? | H_0)P(u_\beta=0 | H_1)]}{\Delta J_0 P(u_\beta=? | H_0) + \Delta J_1 P(u_\beta=? | H_1)} \quad (62)$$

$$W^4 = \frac{\Delta J_0 \Delta J_1 [P(u_\beta=0 | H_0) - P(u_\beta=0 | H_1)]}{\Delta J_0 + \Delta J_1} \quad (63)$$

$$\theta_{2,1} = \frac{P(u_\beta=1 | H_1) \Delta J_1 - C}{P(u_\beta=1 | H_0) \Delta J_0 + C} \quad (64)$$

$$\theta_{2,2} = \frac{[P(u_\beta=1 | H_1) + P(u_\beta=? | H_1)] \Delta J_1 - C}{[P(u_\beta=1 | H_0) + P(u_\beta=? | H_0)] \Delta J_0 + C} \quad (65)$$

$$\theta_{3,1} = \frac{P(u_\beta=0 | H_1) \Delta J_1 + C}{P(u_\beta=0 | H_0) \Delta J_0 - C} \quad (66)$$

$$\theta_{3,2} = \frac{[P(u_\beta=0 | H_1) + P(u_\beta=? | H_1)] \Delta J_1 + C}{[P(u_\beta=0 | H_0) + P(u_\beta=? | H_0)] \Delta J_0 - C} \quad (67)$$

THEOREM 3 : Given the optimum final decision rule u_f of the primary DM (derived in Theorem 1) and the optimum decision rule u_β of the consulting DM, the optimum decision rule for the preliminary decision of the primary DM is given by a deterministic function

$$\gamma_\alpha : Y_\alpha \rightarrow \{0, 1, I\}$$

defined by the following likelihood ratio tests :

$$\gamma_\alpha(y_\alpha) = \begin{cases} 0 & \text{if } \Lambda_\alpha(y_\alpha) \geq \alpha_1 \\ 1 & \text{if } \Lambda_\alpha(y_\alpha) < \alpha_1 \text{ and } \Lambda_\alpha(y_\alpha) \geq \alpha_2 \\ I & \text{if } \Lambda_\alpha(y_\alpha) < \alpha_2 \end{cases} \quad (68)$$

where :

$$\alpha_1 = \begin{cases} \theta_{2,1} & \text{if } 0 \leq C \leq \min\{W^1, W^2\} \\ \theta_{2,2} & \text{if } W^2 < C \leq W^4 \\ \theta_1 & \text{otherwise} \end{cases} \quad (69)$$

$$\alpha_2 = \begin{cases} \theta_{3,1} & \text{if } 0 \leq C \leq \min\{W^3, W^4\} \\ \theta_{3,2} & \text{if } W^3 < C \leq W^1 \\ \theta_1 & \text{otherwise} \end{cases} \quad (70)$$

Remark : (34) and (35) are the equations of the corresponding thresholds for the Gaussian example.

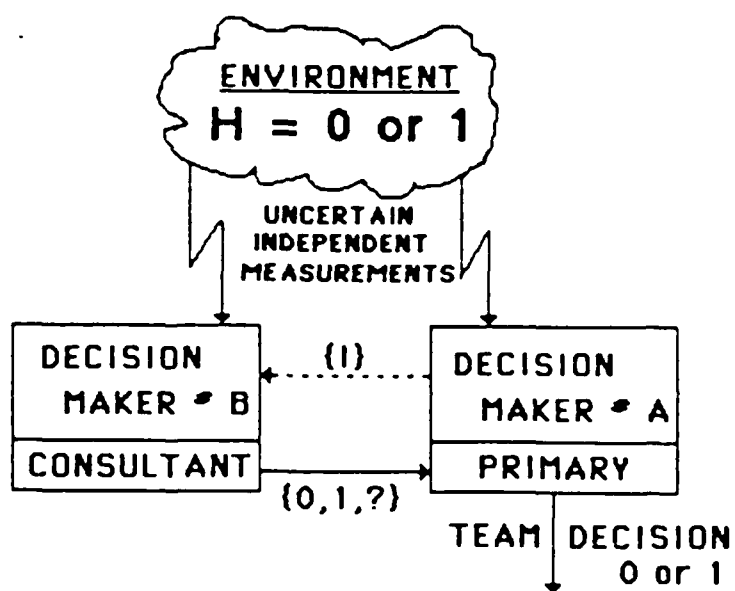


Figure 1. Problem Formulation

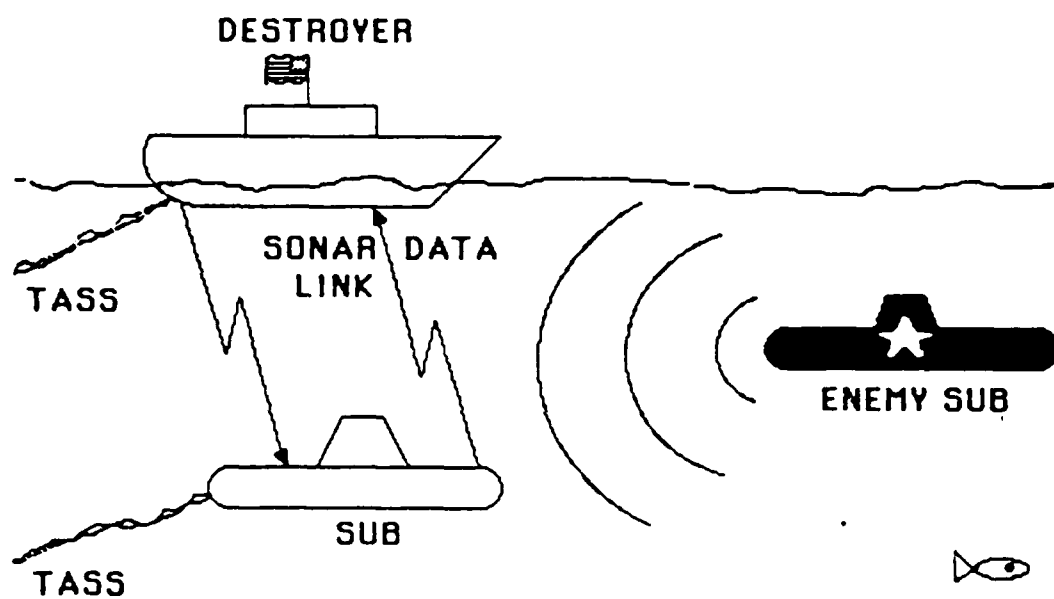
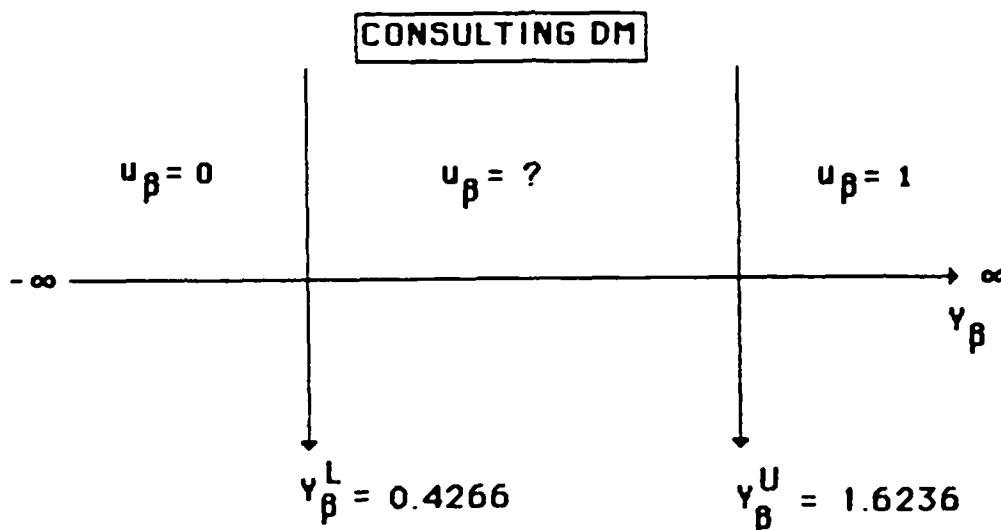
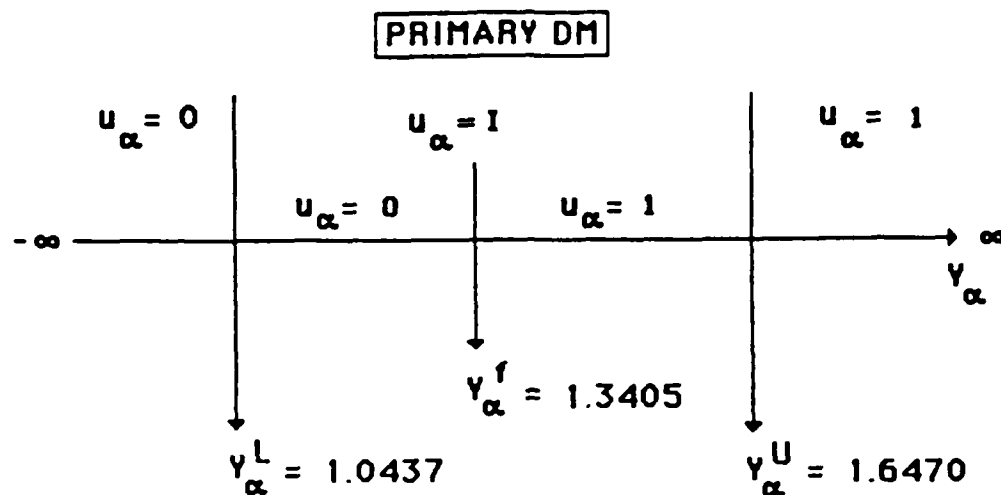


Figure 2. Anti-Submarine Warfare (ASW) Example

* OPTIMAL POLICIES ARE DEFINED BY THRESHOLDS



* BASELINE PARAMETERS

$$\begin{aligned}
 \mu_0 &= -1 & ; & & \mu_1 &= 3 \\
 \sigma_\alpha^2 &= 1 & ; & & \sigma_\beta^2 &= 4 \\
 P(H_0) &= 0.8 & ; & & P(H_1) &= 0.2 \\
 C &= 0.1
 \end{aligned}$$

Figure 3. The Gaussian Case

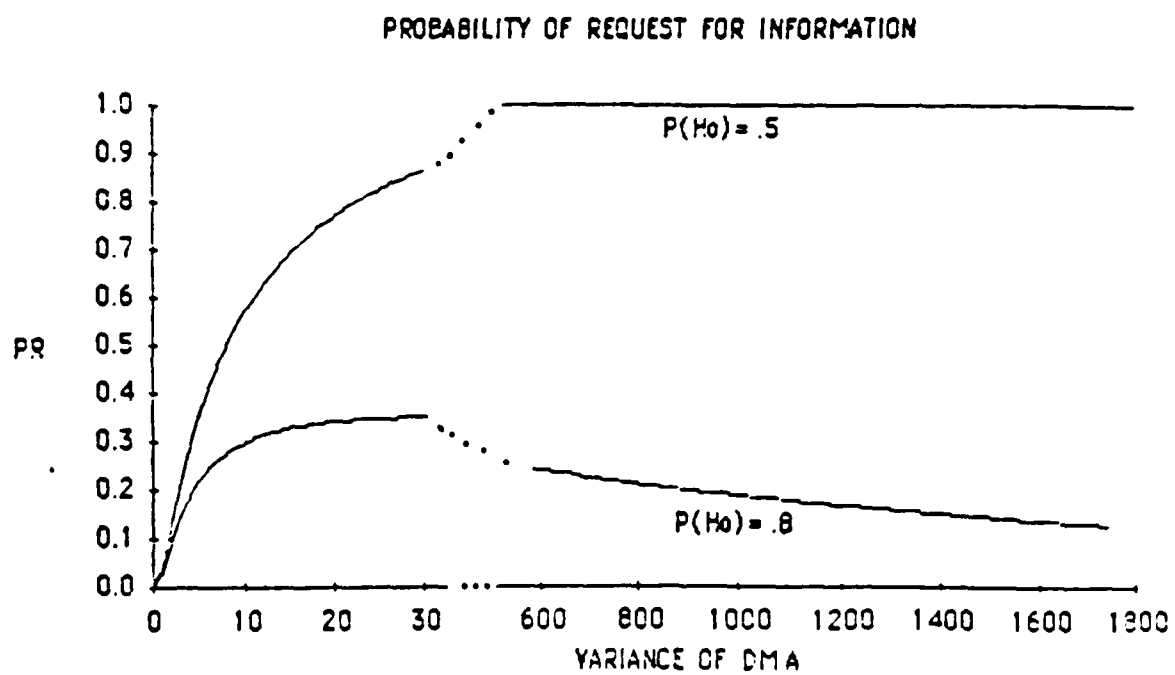


FIGURE 4:

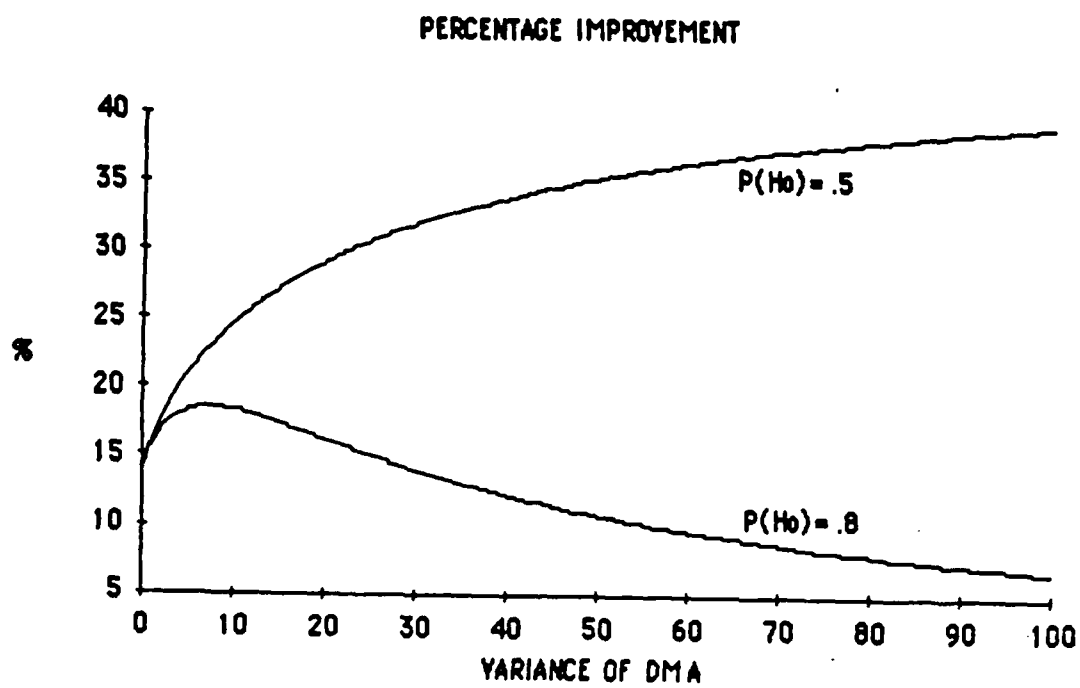


FIGURE 5

THRESHOLDS OF DM A

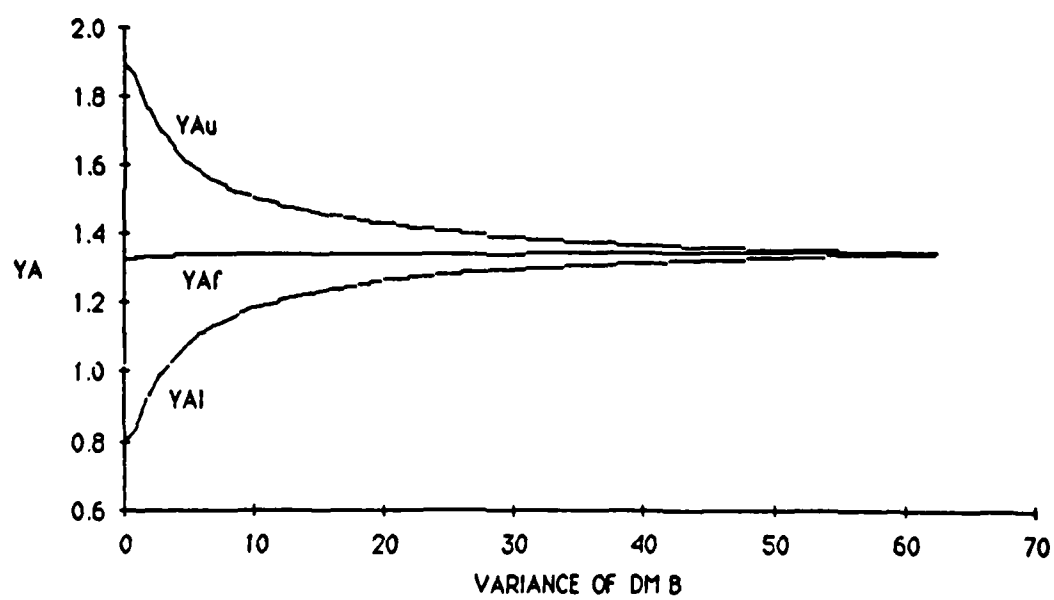


FIGURE 6

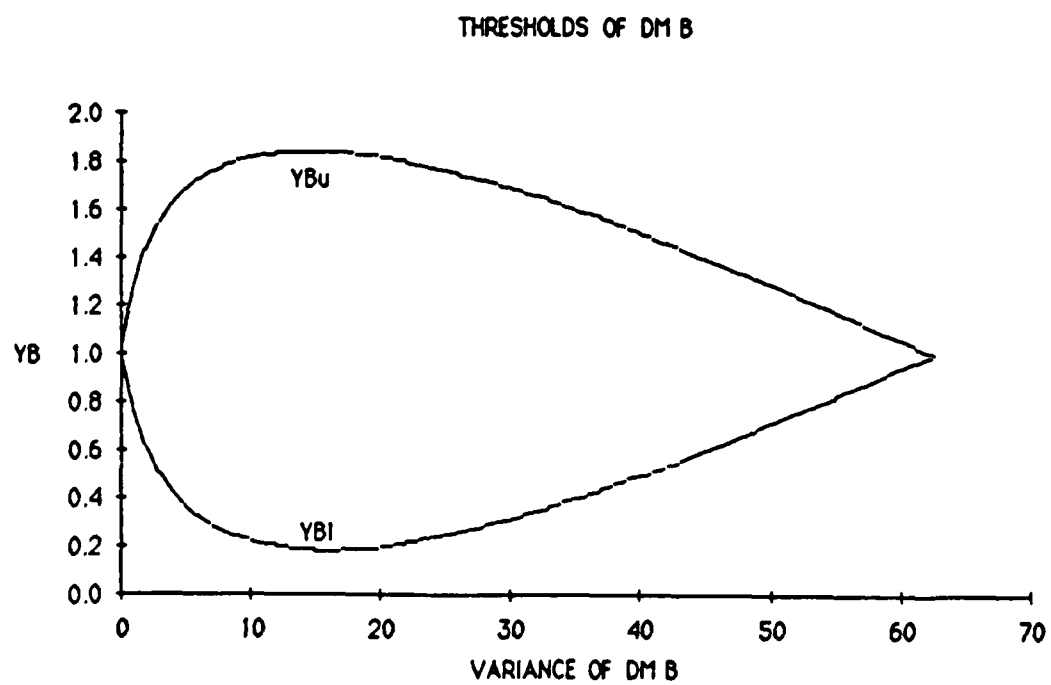


FIGURE 7

THRESHOLDS OF DM B

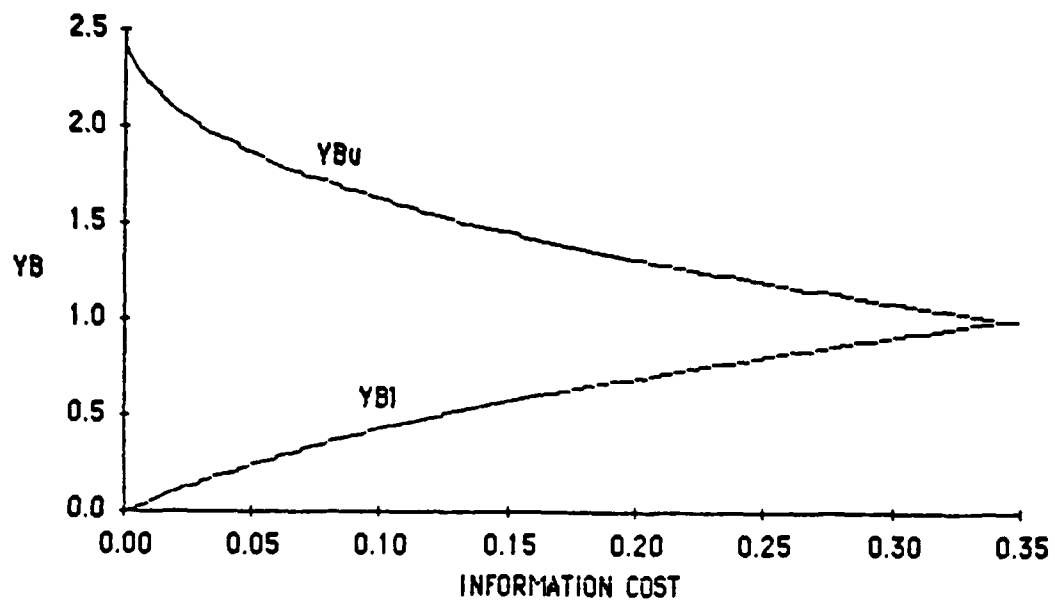


FIGURE 8

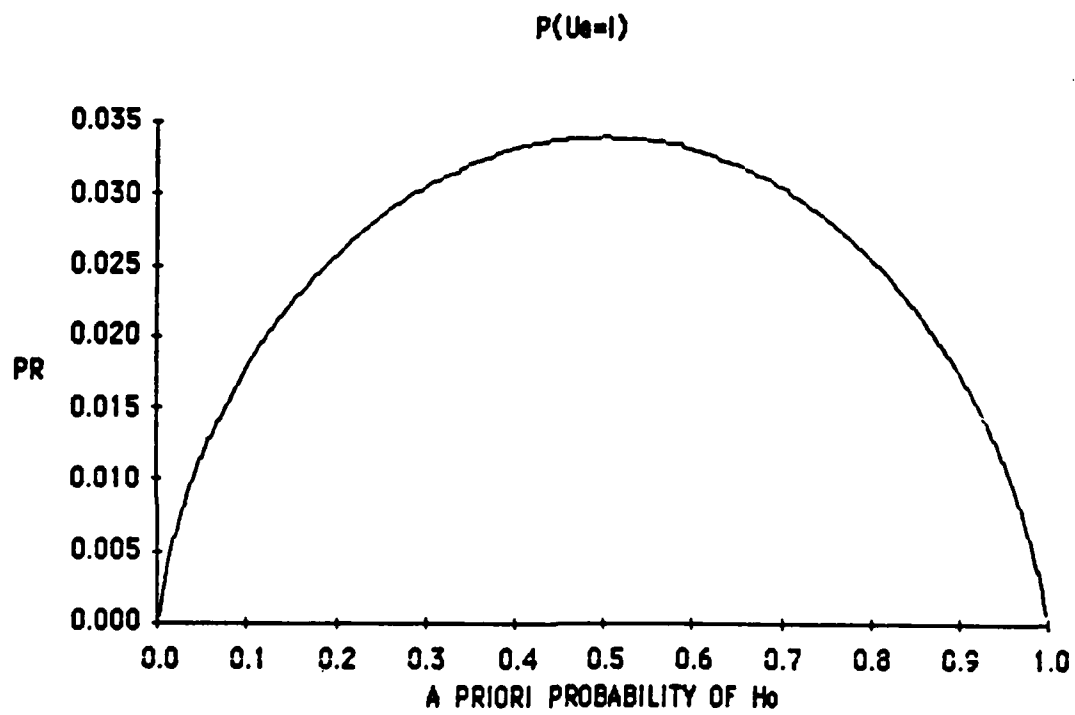


FIGURE 9

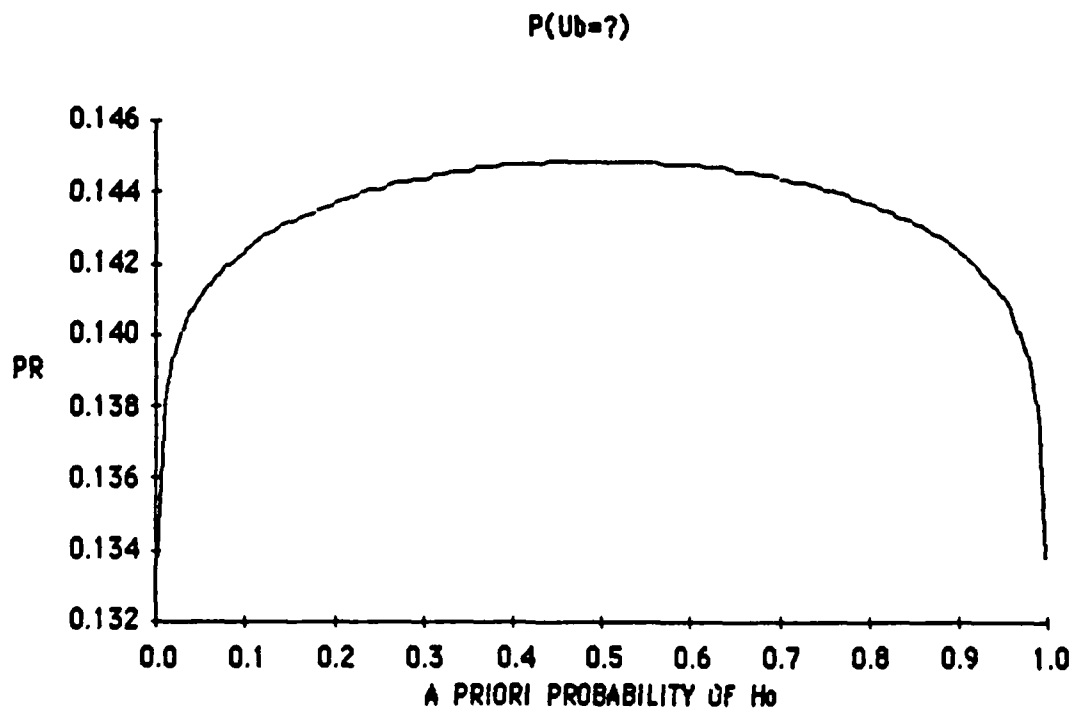


FIGURE 10

PERCENTAGE LOSS IN COST

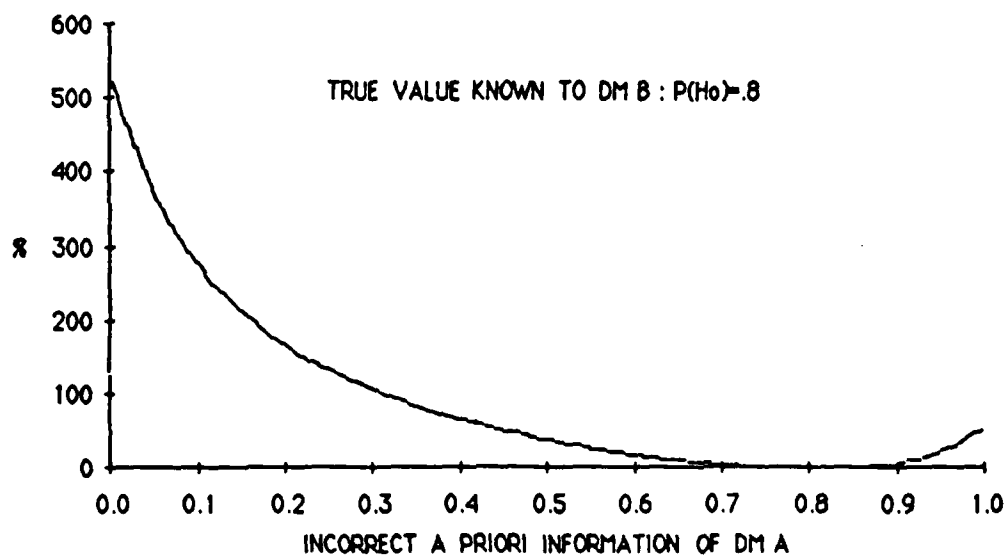


FIGURE 11

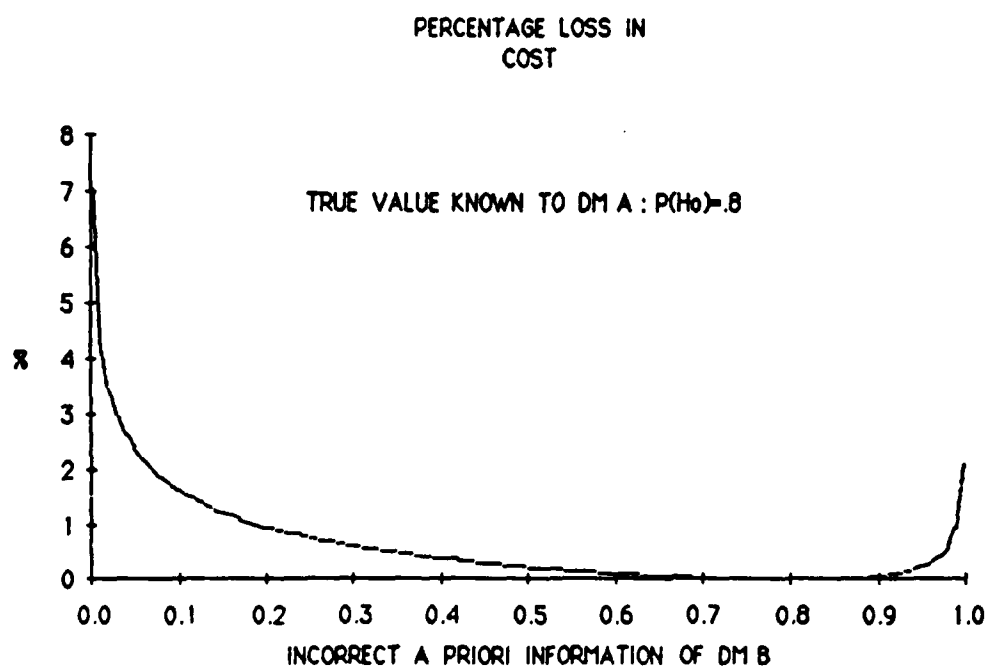


FIGURE 12

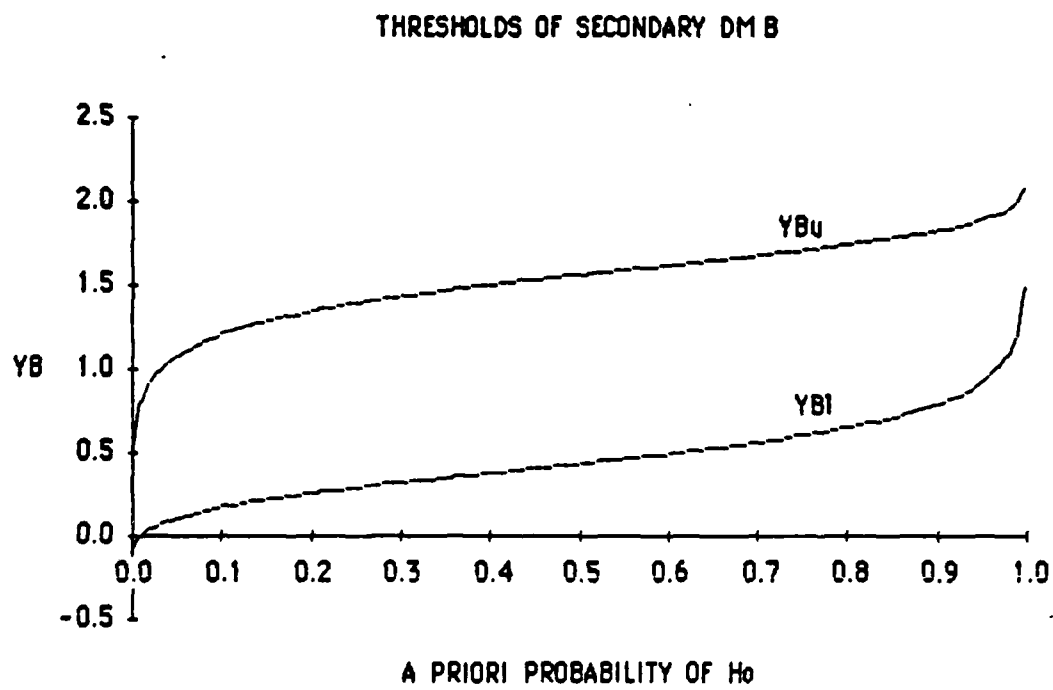


FIGURE 13

END
FILMED

5-86

DTIC